Voice over IP Basics

Net Telephony For All

Before the dotcom bubble burst, Internet and Intranet telephony was regarded as cool technology. Unfortunately VoIP was an inconvenient technology at the time, and it failed to gain a foothold. Now thanks to cheap broadband connections, the time has come to see if we can make use of the technology and save money while doing so.

BY JÖRG REITTER

www.photocase.de

Today, both private users and enterprises can leverage the power of mature VoIP software, and specialist providers offer low-budget connections between the VoIP world and public switched telephone (PST) networks.

Voice over IP (VoIP) is slowly gaining a foothold as a replacement for traditional bell wired networks used by today's analog or ISDN phone systems. VoIP converts voice into IP packets, which can be transported across the Internet at a fraction of the cost. The traditional wiring can remain in place, but changes to the access technology are essential, as VoIP only requires an Internet connection. No more phone bills!

VoIP telephony is quite cheap for private users who simply need a single line. The recommended phone device in this case is a softphone for your Linux PC. You will need a sound card and a headset, but that's really it, as far as capital outlay is involved. If the thought of using your computer as a phone bothers you, consider purchasing a hardware VoIP. VoIP phones are quite expensive at around $150 to $250, but you might be lucky enough to find a VoIP provider who will sponsor your hardware investment. Using a particular VoIP phone supplied by your provider also guarantees that the phone will support the technology required to use that provider's services.

Globally Accessible with VoIP Providers

If the thought of exchanging email messages to organize VoIP phone sessions bothers you, and if you would also like callers to be able to contact you from PST and cellular networks, a VoIP provider is your only option. The provider's job is to assign genuine phone numbers to its customers. On the upside, calls within the provider's network are free. However, you will need a DSL or cable modem connection, as a call with a bandwidth of 80KBits/s up and downstream is more than a single ISDN BRI channel (64KBit/s) can handle.

At this time of writing, there are a limited number of providers that require only a data connection and are thus capable of offering a full replacement for traditional telephony. All of these providers use the open SIP [1] standard and supply the VoIP hardware, such as terminal adapters or SIP hardware telephones either free of charge – or at a surcharge.

Before plunging into the brave new VoIP world, you should compare the prices in Table 1. Some levy a monthly basic charge that includes a certain amount of airtime. Others, do not have a basic charge, but they do ask you to pay in advance.

Providers with a larger number of gateways to the PST network can offer lower rates for calls to external networks, while others are slightly more expensive. The DSL provider charges about 2.5 cents per minute, as it needs to route many calls across the PST network.

Changing Markets

VoIP is very much an emerging market, so customers should compare prices regularly and also analyze their own call patterns. One factor to consider is the fact that not every country is part of the

VoIP providers network, and that calls to other VoIP providers are not free of charge at present. The reason for this is that each provider has its own registrar, and that there is no exchange of data between registrars. The ENUM global messaging directory service looks set to change this.

ENUM as a Global VoIP Directory

The idea of Unified Messaging is to gather all kinds of communication types under a common umbrella: telephony, email, cellular networks and the Web. Addresses of all types will be interconnected transparently for the user, who will need only one application. Users simply register a traditional phone number as a domain. ENUM domains (tElephone NUmber Mapping, RFC 2916, [2]) are DNS-based. This allows for a global directory where phone numbers are linked to a list of URIs [3].

The central e164.arpa registry follows ITU E.164 naming conventions, which also provide the schema for traditional phone number allocations. ENUM is still being tested at present, so you can register a phone number domain free of charge, and data validation is quite quick and safe.

What is special about telephone domains is that they can interconnect various messaging systems, such as IP phones, email, or ISDN phones. Users simply mark checkboxes to specify the order in which to process their URIs when they try to contact someone in the directory. If a caller fails to reach whoever he is calling using the phone, he could simply send an email message for example.

Because ENUM has not expanded as expected, and because of some doubts about the validation methods, another directory was launched in the *e164.org* domain.

This directory uses a callback PIN approach to validation. Users are assigned a PIN for the registered number, and are required to enter the PIN at a specific website. The trouble with e164.org is that you need your own VoIP phone switchboard, like Asterisk, to use the directory service, whereas the ENUM project supports more or less any VoIP provider.

Linux as a VoIP Server

The good news is that there are any number of software phone systems that run on Linux. This is probably due to Linux's reliability and flexibility. The packages are at various stages of development. Some software phone systems, like Asterisk are Open Source, although developed by a single enterprise. Asterisk supports H.323, SIP and the proprietary IAX protocol. It has the reputation of being one of the most complete software phone systems, despite its complexity.

Vovida Networks Vocal [4] VoIP framework follows the same approach. The current 1.5.0 release contains various SIP servers and has a SIP to H.323, a SIP to MGCP gateway, a client and numerous tools in the CVS. The interesting thing about the new version is that it supports Transport Layer Security (TLS) for provisioning The SIP Express Router SER [5] is another SIP-based system that can be used as a registry, a redirection server, or proxy.

Many projects, but so little time

Other programs, such as PBX4Linux by Andreas Eversberg [6] can also hold their sway with these distributions. PBX4Linux only implements H.323 for VoIP at present, but SIP is right at the top of the to-do list. PBX4Linux even has special ISDN adapters (HFC cards) that can replace a hardware based ISDN phone system, as you can connect ISDN phones to them directly. There are more OSS projects, such as GNU Bayonne [7], OpenPBX by VoiceTronics [8], or the Telos ISDN2H323 gateway [9].

Thus, it is not difficult to find free software for a VoIP server system. However, installing the system, and possibly connecting it up to multiple LANs can be a lot of work for admins. Of course, there is the double load of maintaining both systems until the traditional phone system finally goes offline.

Despite the amount of work involved in setting up a Linux VoIP server, just consider the benefits. The system can be used for training purposes, as a backup for a commercial VoIP gateway, or just to cut costs. It pays to keep a level head when considering the costs. First of all, the additional hardware expenditure can be substantial, and availability is still an issue with a combination of Ethernet and IP.

VoIP for Enterprises

Availability is one of the most important issues in telecommunications. If a phone system goes down for a protracted period, loss of face will be the least of your worries. Losing the phone system in a business means losing revenue. In the worst case, it can mean losing several days worth of revenue – a terrible thought!

IP based data networks are not exactly famous for their reliability, in contrast to telecommunications networks like the

Table 1: VoIP Providers

Suppliers	Nikotel (USA, Europe)	Gossiptel (UK)	Vonage (USA/CA)	Freshtel (AU)
Web	*www.nikotel.de*	*www.gossiptel.com*	*www.vonage.com*	*www.freshtel.net*
Cost				
Basic charge per month	6.99 Euro	Free	$14.99	Free
Invoicing rate: First/Following	60/60s	60/60s	60/60s	60/60s
Charge per minute in cents (excerpt)				
Provider network	Free	Free	Free	Free
Home country: PST/cellphone	1.9c/22.7c	2.5p/16.9p	3.9c*/Unknown	4.9c/30c
Transatlantic	2.9c	2.5p	3c	3.9c
Singapore	2.9c	2.5p	5c	3.5c
Connection				
Protocol	SIP	SIP	SIP	SIP
Ability to connect hardware phone	Yes	Yes	Yes	Yes
Hardware provided	–	Phone adapter	Phone adapter	–

* After first free 500 minutes

ISDN PST network. What this means for VoIP is that you not only need a reliable provider, it also shifts the burden of responsibility for telecommunications to the LAN administrator.

Hardware-based phone systems are typically maintained by external service engineers, as the systems are not exactly easy to handle. VoIP places this burden on the admin's shoulders; the admin not only has to deal with faulty hardware, but shoulder the responsibility for software downtime.

Test until you are sure, then test again

VoIP communication is fraught with pitfalls for enterprises, and this makes thorough testing, while the traditional hardware is still running, extremely important.

The easiest way to do this is to use softphones to connect people across your LAN. This setup has a few advantages: First of all, the business does not need a provider; it can avoid NAT or firewall configuration issues, and take care of VoIP teething trouble, without impacting the business. This setup simply assumes that the VoIP system is attached to the LAN. After successfully completing testing, the VoIP system can be linked up to the existing phone system to allow for a gradual migration (see Figure 1).

VoIP adds very little traffic to a local network. Depending on the codec, a call can use between 6 and 64 KBit/s of netto bandwidth. A 100 MBit Ethernet network should be able to provide this capacity without any trouble at all. Check out the list of codecs at [10].

If you have a small network with less than 50 nodes, a Pentium or Athlon machine should be sufficient for a software based phone system. Your budget planning should leave some leeway for headsets and sound cards for sound output.

Hardphones and the Electricity Bill

If you pull the plug on your ISDN phone, you will note that it only has one connector. The phone uses this connector to attach to the bell wire. The good thing about analog or ISDN phones is that the bell wire also provides the power. Base stations for mobile phones, such as DECT phones, need a power supply, but DECT is not typical in business environments.

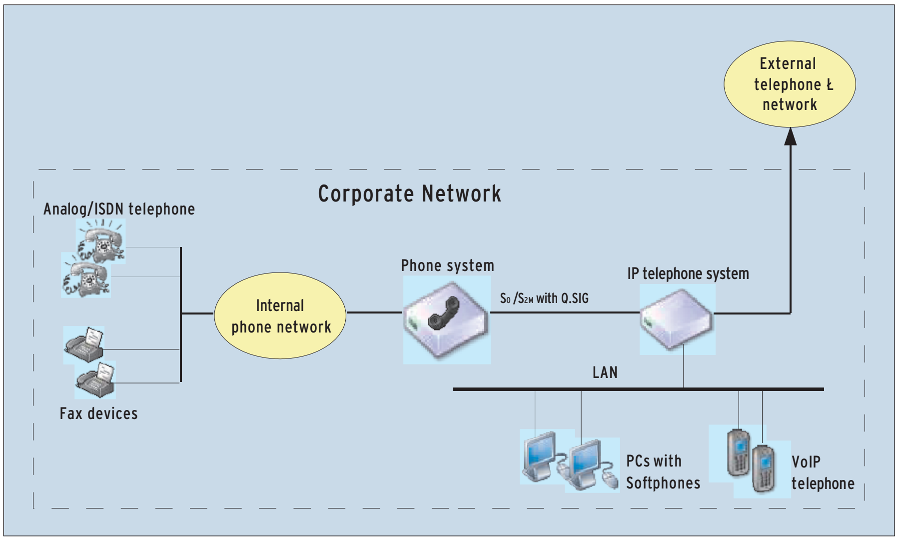

Figure 1: VoIP traffic from the LAN and calls from the phone system interface at the VoIP system. At the same time, the IP system connects the enterprise to the PST network.

VoIP phones follow a similar pattern. Again there are two types of hardphone. One of these allows users lots of mobility, and uses a WLAN connection to link the receiver to the base station. The other type is powered via the network cable. This approach is known as Power over Ethernet (PoE) and was standardized as IEEE 802.11af in 2003.

PoE is a practical thing – assuming that your infrastructure is based on copper wire. Besides the telephones, you simply need new switches that supply node devices with power and network traffic.

Enterprises with fiber to the desk can rule out this low-budget approach, and will need to install either hybrid cabling (which could have a nasty impact on their budgets), or attach the phones to the power supply, which could produce a bird's nest of power cabling.

If you need to upgrade the infrastructure, you should consider opting for state-of-the-art copper cable rather than FDDI. Gigabit Ethernet on copper requires CAT 6 or 7 wires. CAT 6 does not leave a lot of leeway, with a bandwidth of 250 MHz; thus the 600MHz CAT 7 variant is probably the better choice. CAT 8, which will have a nominal bandwidth of 1200MHz, and thus be capable of transporting digital video data, is another candidate, although not standardized at present.

An Overview of the VoIP Protocols

Protocols in a VoIP network handle the setting up and terminating of connections, signaling, or opening channels for multiple extension calls, such as conferences.

The International Telecommunications Union ITU introduced the first VoIP protocol, H.323 [11], in 1995. H.323 is closely modeled on ISDN and applies the ISDN schema to the world of IP.

In contrast to this, the Session Initiation Protocol, SIP, which the IETF (Internet Engineering Task Force) specified in RFC 3261 [12], is modeled on protocols like HTTP. This leaves developers plenty of room to experiment and find the right balance. Both SIP and H.323 define distributed architectures and are peer-to-peer protocols.

The H.248/Megaco extension of the Media Gateway Control Protocols MGCP [12] provides a centralized solution for creating multimedia and voice applications. The extension was a joint effort between the ITU and the IETF. H.248/Megaco works in parallel with H.323 and SIP, separating the load from the signals, and is only used at network borders, where call agents communicate with unintelligent gateways that simply forward the traffic load.

H.323 Was There First

As an early development, H.323 is closely modeled on the H.320 recommendation that specifies the video-phone, and thus combines voice and video. Many channels have been introduced over the years, and last year saw the release of version 5 of this protocol.

Numerous manufacturers pounced on H.323 and integrated the protocol into their VoIP products. The reason why the standard was so readily accepted is that it stipulates more or less everything. But there is a growing tendency for manufacturers to ditch the specs. Today H.323 suffers from an overdose of incompatible terminal devices and gateways. IT managers are thus at the mercy of their hardware suppliers.

The fact that H.323 uses a number of really expensive protocols is another bad thing for newbie developers. This proliferation of protocols is a thorn in the side for service providers, as setting up a H.323 network is extremely complex, and that makes the setup all the more error-prone. Converting signals to binary format was an approach that was welcomed at first, although it merely added another level of complexity, as it turned out.

Although it allows the protocol to be more conservative in its bandwidth usage, the effort is a little self-defeating, as the conversion process needs more CPU power and takes far longer to complete.

Good News is around the corner

Now for the good news about H.323. Being the most venerable of all VoIP protocols, many devices support it, and that means a more competitive market. The ITU recognized the value of distributing the additional load sensibly across the network, and introduced integrated load balancing at an early stage. Also, H.245 signaling supports encryption. If you are interested in VoIP development, check out the OpenH323 [13] website. The project provides an Open Source library released under the Mozilla Public License for programmers.

SIP – the new Challenger

Shortly after SIP entered the VoIP stage, a debate on the relative merits of the protocols broke out. The only thing one can say for sure at present is that H.323 devices are more widespread, whereas SIP is more popular with manufacturers on account of its characteristics. Many Linux PBX systems or softphones support SIP.

SIP is based on well-known Internet protocols such as HTTP, and is defined as an Application Layer protocol, where placing and terminating VoIP calls, event notification, or multimedia sessions are typical applications. Multimedia conferences, and VoIP calls that use voice and video, E-learning or the like are typical multimedia applications. oSIP by the GNU project [14] is one example of an Open Source implementation.

Besides HTTP, SIP also uses some other IETF protocols. DNS (Domain Name System) and URLs are used to handle naming. Nodes use the Session Description Protocol (SDP) to exchange capabilities, and Multipurpose Internet Mail Extensions (MIME) allow for application integration. This makes it easy to integrate SIP with other Internet technologies and/or applications. The addressing schema should be familiar to you: Users are assigned addresses in email format, such as *sip:anton@foo.com*. RFC 2806 adds the aspect of so-called Tel-URIs, such as *tel: +99.12345.6789*, which are modeled on normal phone numbers.

SIP is quite flexible with respect to its choice of transport protocol. It can use UDP, TCP, SCTP. UDP is the most efficient protocol, as it has the least overhead. The fact that SIP can handle retransmitting itself additionally simplifies the handling of firewalls and multicast applications. On the downside, SIP does not support resource reserva-

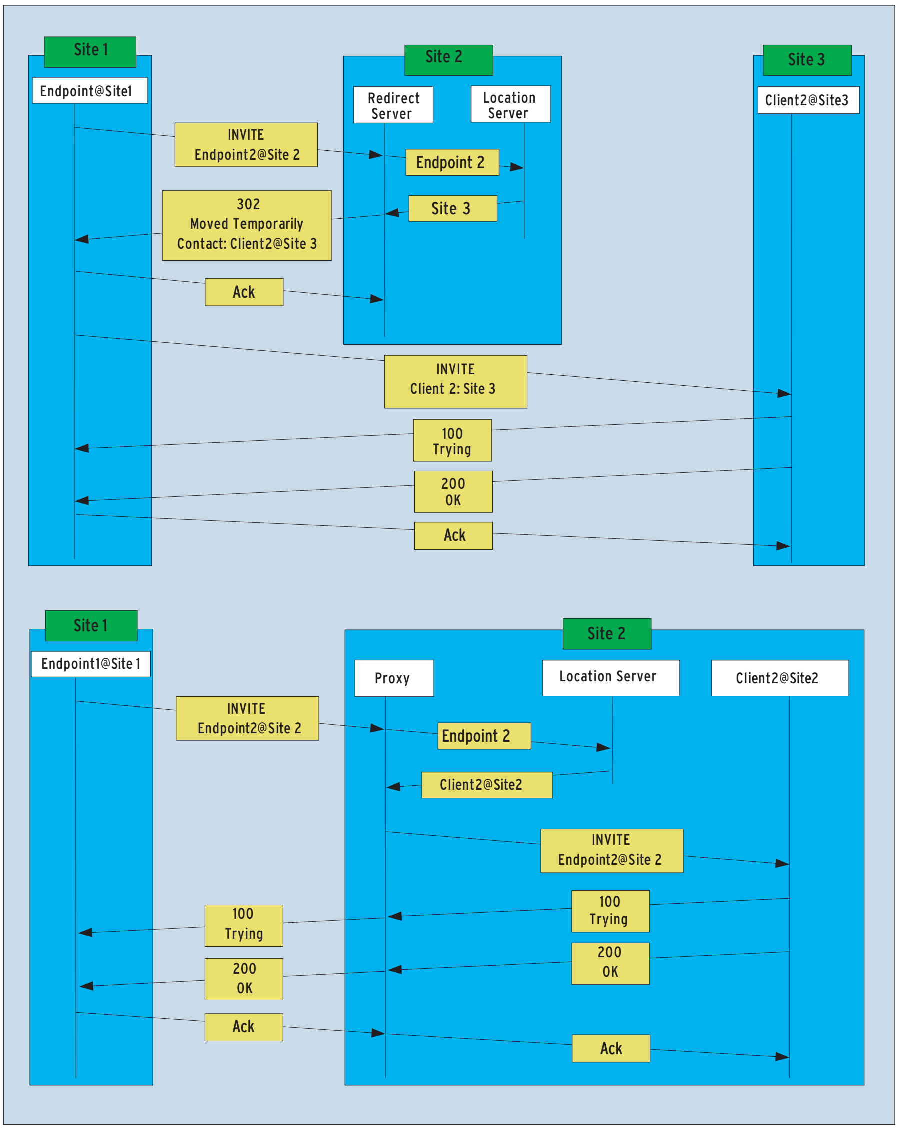

Figure 2: The path a SIP call takes through the SIP network, showing redirect mode (top) and proxy mode (bottom).

tions in a basic setup. However, nodes can use SDP to exchange service classes, although this requires all the devices on a specific path to support the same class.

SIP a Detailed Look

A SIP network comprises of multiple software components that handle different tasks. An SIP User Agent (a IP phone, SIP gateway, or PC) initiates and terminates a session. A SIP registry server maps names to addresses in a similar way as to a Domain Name Server. The SIP network also defines two proxy types (stateless and forking), and a redirect server which accepts SIP requests and forwards them on to a new location when necessary.

Common Language

All of these elements speak a common request/response language. The basic specification distinguishes between the methods *Invite*, *Ack*, *Options*, *Cancel*, *Bye* and *Register*. When a caller *invite*s another user, the user agent transmits a SDP session description. This can describe a one on one call, or a conference. The user agent that receives the call acknowledges (*Ack*). And *Bye* terminates the call. Figure 2 shows the handshake procedure.

User agents communicate by means of a Real Time Protocol (RTP) stream. SIP transactions are handled by proxy servers which use the options method to exchange the functions supported by the agents and servers. The cancel method interrupts the current session. The registrar method registers a client with the location service.

Responses are numerical codes that indicate the success or failure of a request, and additionally specify a redirection address, if the receiver of the call happens to be elsewhere than originally expected.

Thus, a request to *sip:anton@foo.com* might be routed to the office first. If noone answers, the forking proxy will route the request to the user's cellphone, and then to an alternative address or a URL. This process is continued until a positive response is finally returned to the caller, ensuring the call is not lost.

Proxies forward requests for user agents, using DNS servers for addressing purposes. A proxy learns the whereabouts of a user agent from location servers, which they query for call routing information. Users can update the user databases on the location servers by means of their agents, which register with the servers.

SIP and (In-)security

IT managers that deploy an open Internet architecture such as SIP can expect security issues. Just like in traditional IP networks, there are internal and external sources of danger. Attackers are difficult to locate and trace, and there are many attack vectors. An internal attacker could spoof packets. Calls can be redirected by manipulating the from entry in the register method. A similar entry in the invite entry might avoid a call filter.

Also, attackers can sniff packets off the wire or signaling path. If an attacker can grab RTP packets, she might even be able to eavesdrop on conversations. If an attacker can sniff SIP packets, she might be able to decode communication links. Spoofed packet IDs could be used to evade billing mechanisms, and there is also a danger of replay attacks that use packets grabbed off the wire.

Finally, Denial of Service (DoS) attacks might completely cut an enterprise off from the outside world. To make DoS attacks more difficult, administrators can harden their SIP device configurations and transmission elements such as proxies. In addition, network based Intrusion Detection or Intrusion Prevention systems should be used to monitor suspicious activity.

An encrypted transport layer would prevent attackers from sniffing packets. IPsec or Transport Layer Security (TLS) are both useful, and provide far more security, despite being relatively complex, as they need to be implemented throughout the network to be effective. Media datastreams can also be secured using Secure RTP.

Admins can use authentication technologies to prevent spoofing attacks. SIP Invite/Register methods between user agents and proxies can be secured using HTTP digest authentication or certificates.

The former uses a challenge-response method based on a pre-shared key. The certificate-based approach involves signing parts of the header and the request. The third approach, HTTP Basic Authentication, uses a password, but unfortunately for most basic security purposes, one that is transmitted in the clear.

Testing VoIP Without Risk

Users wanting to migrate to VoIP are lucky. The low prices that today's VoIP providers offer auger well in view of an increasingly competitive market. The ENUM domain will provide a globally accepted directory that unifies modern telecommunication media, and allows users to quickly locate and contact whoever they are looking for.

The Open Source Community has a good selection of VoIP clients for H.323 or SIP. The clients are often available both for Linux and Windows; and this is an advantage in heterogeneous environments (or on dual-boot machines). Server-side there are a number of soft PBX systems for Linux.

They connect the ISDN telecommunications system with the VoIP network, without endangering your past investments. On the contrary, they extend the capabilities of the existing system. Also the fact that the technology is license free, opens up the way for testing and evaluation. ■

INFO

[1] SIP: *http://www.ietf.org/rfc/rfc3261.txt*

[2] ENUM-RFC: *http://www.ietf.org/rfc/rfc2916.txt*

[3] ENUM: *http://www.enum.org*

[4] Vocal: *http://www.vovida.org/applications/downloads/vocal/*

[5] SIP Express Router: *http://iptel.org/ser/*

[6] PBX4Linux: *http://isdn.jolly.de/*

[7] GNU Bayonne: *http://www.gnu.org/software/bayonne*

[8] OpenPBX: *http://www.voicetronix.com/open-source.htm*

[9] ISDN2H323: *http://www.telos.de/linux/H323/default_e.htm*

[10] Codecs: *http://www.voip-info.org*

[11] H.323: *http://www.packetizer.com/iptel/h323/*

[12] H.248/Megaco: *http://www.faqs.org/rfcs/rfc3015.html*

[13] OpenH323: *http://www.openh323.org/*

[14] oSIP: *http://www.fsf.org/software/osip/osip.html*